

# Leveraging the Web for Services at



Mark Nottingham  
<mnot@yahoo-inc.com>



1992 1994 1996 1998 2000 2002 2004 2006 2008

100,000  
users

25 million  
users

200 million  
users

500 million  
users


49 employees

11,400 employees

65 million  
daily page views

1.3 billion  
daily page views

4 billion  
daily page views



# Integration Nightmare

Y! Auctions

Y! Mexico

Y! Messenger

Y! India

**Bix**

Y! Entertainment

Y! Travel

Y! China

Y! Germany

**Inktomi**

**Upcoming**

Y! Sports

Y! Korea

Y! Tech

Y! Small Business

Y! Brazil

**del.icio.us**

Y! Mail

Y! Food

Y! Japan

**Flickr**

Y! UK

Y! Go

Y!

Y! 360

Y! Spain

**Overture**

Y! France

Y! Health

Y! Music

Y! Italy

Y! Search

Y! Australia

Y! Argentina

Y! Chinese

Yahooligans!

Broadcast.com

Y! Games

**LAUNCH**

Y! Video

Y! Photos

**Kelkoo**

Y! Singapore

**GeoCities**

Y! Maps

Y! Classifieds

**HotJobs**

Y! Podcasts

Y! Shopping

Y! Groups

Between  properties

With  partners

With  acquisitions

Within  over time

“!™ Scale”



Web | Images | Video | Local | Shopping | more

Search:  Web Search

Y! Answers: Ask a question | Answer questions

the Web  UK only

My Yahoo!

My Mail

Page Options

- Autos
- Finance
- Games
- GeoCities
- Groups
- HotJobs UPDATED!
- Maps
- Movies
- Music
- My Web
- Personals
- Photos
- Real Estate
- Shopping
- Sports UPDATED!
- Tech
- Travel
- TV
- Yellow Pages
- Y! Telemundo NEW!

More Yahoo! Services

Small Business

- Get a Web Site
- Domain Names
- Sell Online
- Search Ads

Featured Services

- Downloads
- Health
- Kids
- Mobile
- Voice
- UK Broadband
- Y! International

Featured Entertainment Sports Life

Mar 15, 2007



The eyes of March?

Searchers want the scoop on March 15, but what word are they really looking for?

» Find out

- Check out the drama of TV's 'Rome'
- Y! Answers: How did Caesar die?

Beware the Ides -- or is it 'eyes?' -- of March

Why America embraces March Madness

Some courting birds get help from a 'wingman'

Final 'Potter' set to break publishing records

» More Featured

In the News World Local Video

As of 12:59 p.m.

- Al-Qaida No. 3 says he planned 9/11
- Republican says Gonzales should be fired
- World powers agree on new Iran sanctions
- 49 police officers killed in India
- French Resistance hero Lucie Aubrac dies at 94
- Debate continues over excavating Chinese imperial tombs
- Are consumers ready for an \$11K cup of coffee?
- NCAA Tournament · MLB · NBA · NFL · NHL · Soccer

» More: News | Popular | Business

Markets: Dow: +0.5% Nasdaq: +0.9% Sponsored by: **Scottrade**

Marketplace



Why online degrees are hot:  
1) Top schools 2) Drive own schedule 3) Employer may pay tuition - 4) Earn degree while working. Earn more.

What's your credit score 560? 678? 720? The average U.S. credit score is 675. See yours for \$0. By Experian.

Attractive card for ugly credit. - A MasterCard for those with less than perfect credit.

Countrywide® - Refi to combine your 1st mortgage and other debt - Bad credit? Low refi rates. 4 of 5 approved.

Hi, Mark

Sign Out

Mail	Messenger	Radio
Weather 49°F	Local	Horoscopes

Yahoo! Shopping

» Shop now

Buying Guides for Top Products

MP3 Players	Bed & Bath	Camcorders	Jewelry
Furniture	Laptops	Kitchen	TVs

Inside Yahoo! Answers



Queen Rania wants to know

What are you doing to empower women in your community?

» Answer Now

Pulse - What Yahoos Are Into

Top-Rated Indian Restaurants



- |  |  |
|--|--|
| • <b>Tabla</b><br>New York, NY               | • <b>Tandoor Indian</b><br>Dallas, TX    |
| • <b>India's Oven</b><br>Los Angeles, CA     | • <b>Bombay Mahal</b><br>Boston, MA      |
| • <b>India House</b><br>Chicago, IL          | • <b>Himalayas</b><br>Atlanta, GA        |
| • <b>Vik's Chaat...</b><br>San Francisco, CA | • <b>Gandhi India's</b><br>Las Vegas, NV |

» Find More Curries and Tandooris

Popular Searches: Today's Overall Leaders

1 NCAA Basketball

5 American Idol



# Media Group

Y! Music

Y! TV

Y! Personal Finance

Y! Kids

Y! Finance

Y! Movies

Y! Fantasy Sports

Y! Games

Y! Health

Y! Sports

Y! Astrology

Y! Wii

Y! Tech

Y! Food

Y! News

Sports

Outdoors

Worldwide

Fantasy

Jan 22, 2002

[Top](#) [NFL](#) [NBA](#) [NHL](#) [MLB](#) [NCAA: Football](#) | [Men Hoops](#) | [Women Hoops](#) [Golf](#) [Tennis](#) [Soccer](#) [more](#)

## Yahoo! Sports

## Today

## NBA

[Scoreboard](#) [Courtside Live](#)

<a href="#">Chicago</a> vs. <a href="#">Miami</a>	0	7:30 ET
<a href="#">Atlanta</a> vs. <a href="#">Orlando</a>	0	7:30 ET
<a href="#">Phoenix</a> vs. <a href="#">Milwaukee</a>	0	8:00 ET
<a href="#">Philadelphia</a> vs. <a href="#">Washington</a>	0	8:00 ET
<a href="#">Dallas</a> vs. <a href="#">Houston</a>	0	8:30 ET
<a href="#">New Jersey</a> vs. <a href="#">San Antonio</a>	0	8:30 ET
<a href="#">Denver</a> vs. <a href="#">LALakers</a>	0	10:30 ET
<a href="#">Sacramento</a> vs. <a href="#">Portland</a>	0	10:30 ET

## NHL

[Scoreboard](#)

<a href="#">NY Rangers</a> vs. <a href="#">NY Islanders</a>	0	7:00 ET
<a href="#">Ottawa</a> vs. <a href="#">Philadelphia</a>	0	7:00 ET
<a href="#">Washington</a> vs. <a href="#">Atlanta</a>	0	7:30 ET
<a href="#">Toronto</a> vs. <a href="#">Calgary</a>	0	9:00 ET

[NCAA Hoops \(M\)](#) [Audio](#)  
[Scoreboard](#)

5 <a href="#">Cincinnati</a> vs. <a href="#">St Louis</a>	0	9:30 ET
7 <a href="#">Virginia</a> vs. <a href="#">Georgia Tech</a>	0	7:30 ET
9 <a href="#">Kentucky</a> vs. <a href="#">Auburn</a>	0	9:00 ET

## Gruden packing his bags?

Coach Jon Gruden will not sign a new contract with the Oakland Raiders **and plans to leave the team after next season at the latest**, his agent told several San Francisco Bay area newspapers. Gruden, a top candidate for the Tampa Bay Buccaneers' job, is in the lower half of the NFL coaching pay scale, with an annual salary of \$1.2 million -- and his parents live in Tampa.



Gruden has refused to discuss his future in detail. (AP) [more photos](#)

## Top Stories

- [Seles beats Venus Williams, moves into semifinal against Hingis](#) - Tennis
- [Gruden's agent says coach won't stay in Oakland past 2002](#) - NFL
- [NBA announces new six-year, \\$4.6 billion deals with ESPN-ABC, TNT](#) - NBA
- [Tyson to fight Lewis in April](#) - Boxing
- [Former Winter Olympian Jack Shea dies after car crash](#) - Olympics
- [Burnitz, Rusch and Zeile move in three-team, 11-player trade](#) - MLB
- [Hornets 111, Knicks 68](#) - NBA
- [X-rays show Warner has bruised ribs](#) - NFL
- [Mavericks 113, Nets 105](#) - NBA
- [Duke is unanimous No. 1 again](#) - NCAA Men's Basketball

headlines updated at Tue Jan 22 09:05:51 2002

## Broadcast &amp; Chat Events

## - What's On Today -

- **MLB Chat:** [Rick Helling](#) - 2 pm ET
- **Men's NCAA Hoops:** [Michigan St](#) at [Iowa](#) - 6:15 pm ET
- **NBA:** [Philadelphia](#) at [Washington](#) - 7 pm ET
- **NCAA Hockey:** [Brown](#) at [Providence](#) - 7 pm ET

## Media

- [NBA Highlights](#)
- [View Slideshow](#) NEW!
- [Sports Talk Radio](#)
- [Photo Gallery](#)


## Fantasy Sports

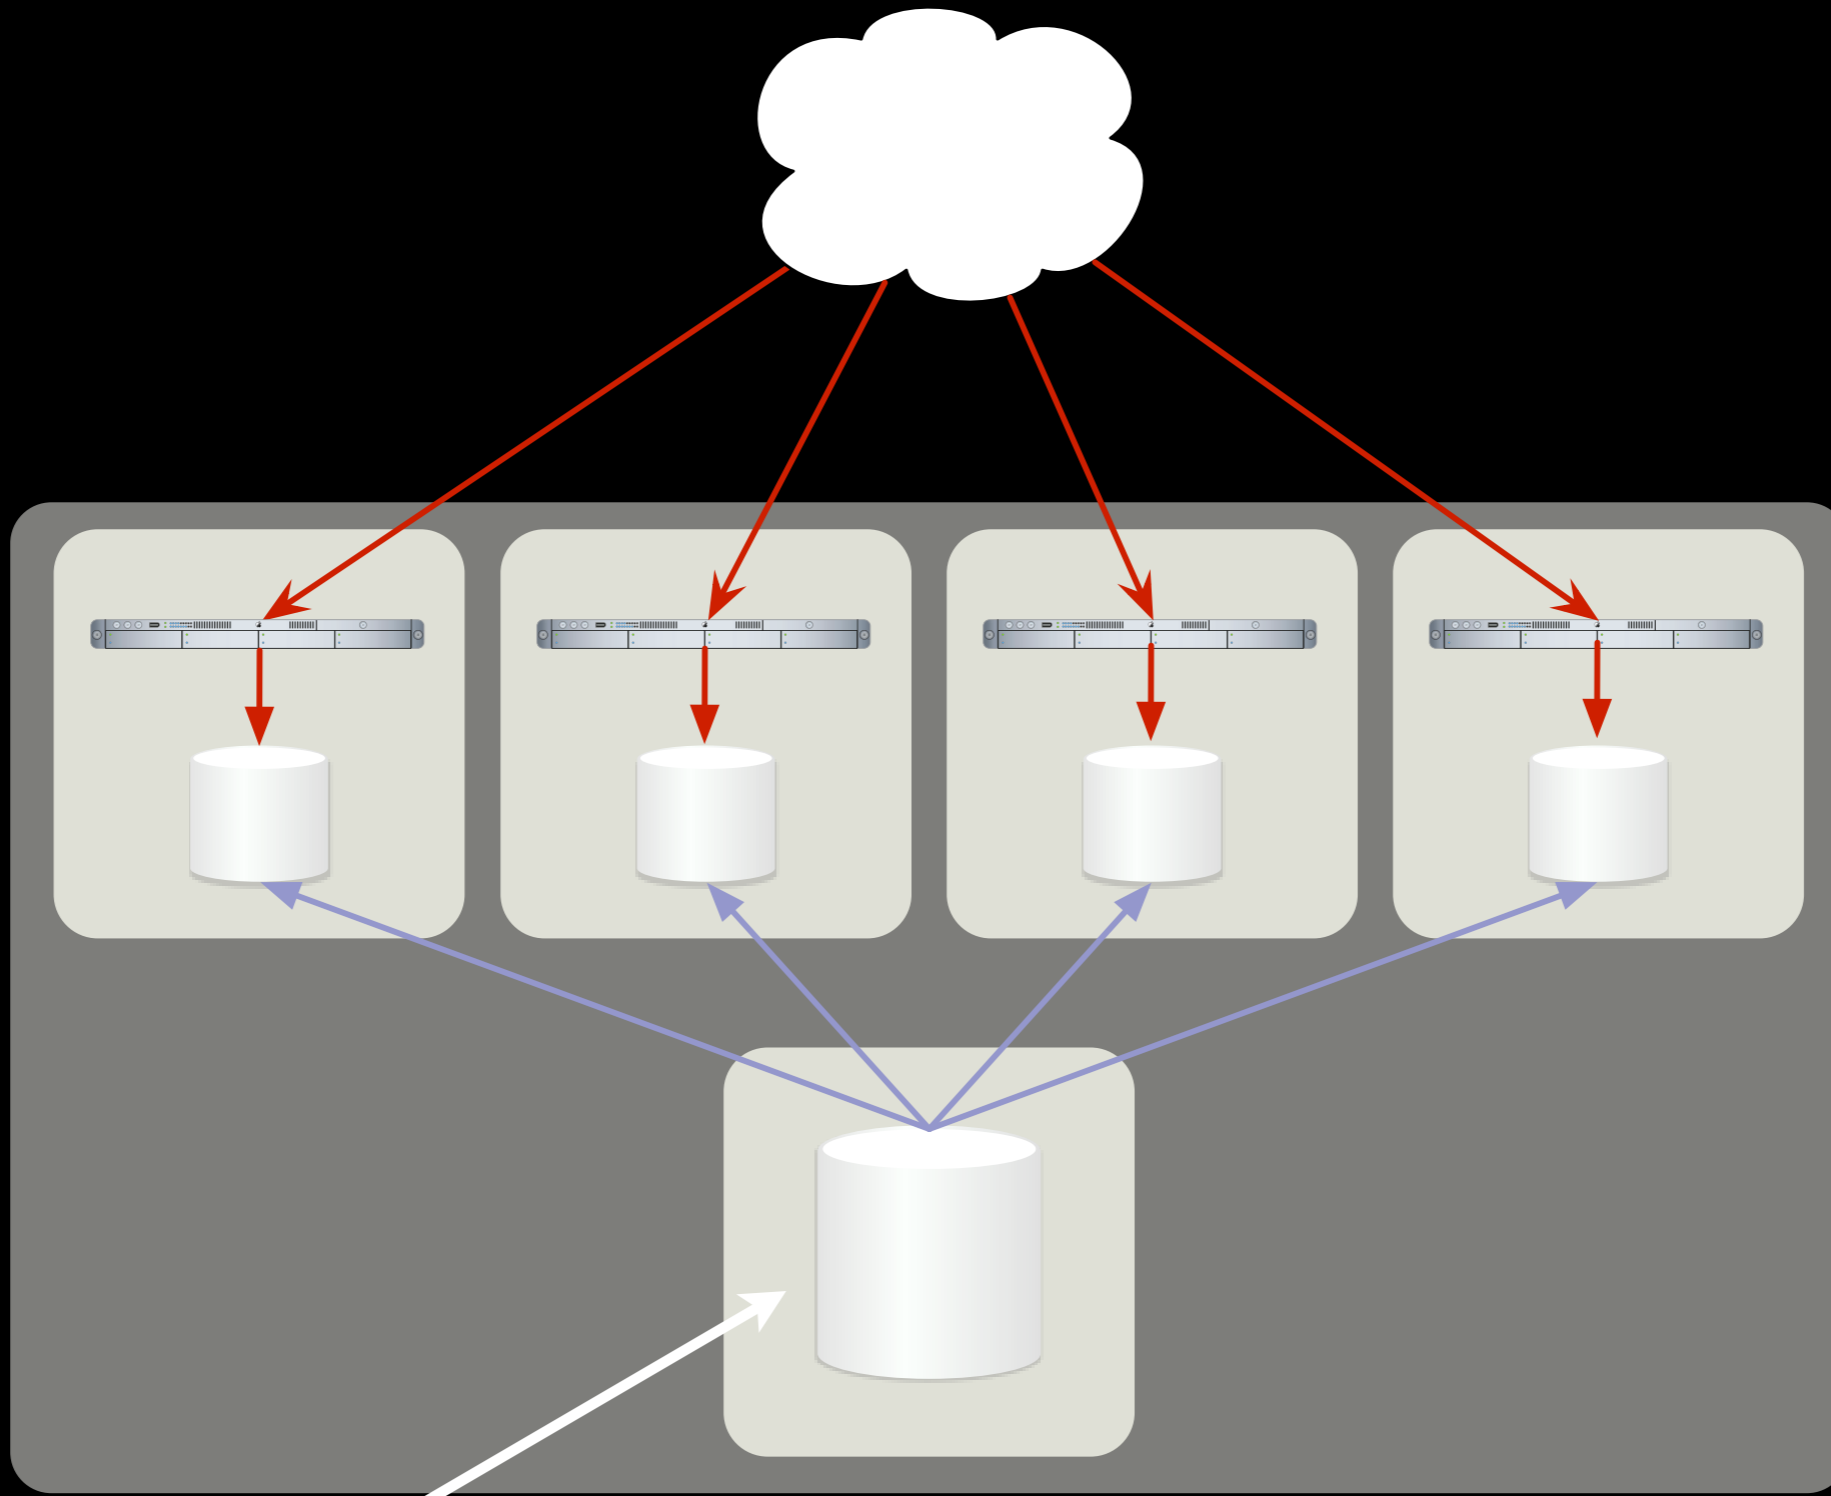
- [Golf](#) | [Auto Racing](#)  
 - [College Hoops](#) | [More](#)

## User Survey

[Help us provide better services for you!](#)








---


Large Datasets Don't Push Well

# Adding Capacity is Expensive



Large Datasets Don't Push Well

# Adding Capacity is Expensive




Large Datasets Don't Push Well

What to do after push failure?

User-Generated Content

Adding Capacity is Expensive



Large Datasets Don't Push Well

What to do after push failure?

Volunteer Abroad with GVI

Join over 60 volunteer abroad programs around the world from 1 to 28 weeks.

www.gviusa.com

HOME

ADVISORS

TECH SHOWS

MY TECH

Search Yahoo! Tech:

SEARCH

### My Tech

#### My Saved Tech Products

Select All

COMPARE PRODUCTS

EMAIL

? What's this?



Remove

Kodak EasyShare C875



Remove

Canon PowerShot SD500 Digital ELPH



Remove

Apple MacBook Pro



Remove

Samsung BlackJack SGH i607

Find all your shopping lists at Yahoo! Shopping.

#### Questions I've Asked



How do I sync a SonyEricsson k610i to a Mac?

No answers yet

To see your Questions and Answers in other categories, visit Yahoo! Answers.

#### Questions I've Answered



Me and my friend have a Sony Ericsson W850i. I am orange he is on t-mobile. why are are menu items differant?

2 Answers

No Products In List

MY TECH



mnotting

MORE MY TECH

My Product Ratings & Reviews

#### My Saved Tech Products

ADD A PRODUCT

ADD A PRODUCT

ADD A PRODUCT

ADD A PRODUCT

ADD A PRODUCT

ADD A PRODUCT

QUICK COMPARE

SHOW ALL

#### Gadgets You Might Like

Why was I recommended these gadgets?



Samsung Trace Cell Phone

USER RATING  
★★★★★



Syntax Olevia 542I Television

USER RATING  
★★★★★



Canon PowerShot A640 Digital Camera

USER RATING

- Sources
  - Yahoo! Search
  - Yahoo! Local
  - Fetch
  - Google Base
  - Flickr
- User inputs
- Operators
- Url
- String
- Date
- My pipes

**Search For (text)**

Name:

Prompt:

Position:

Default:

Debug:

**URLBuilder**

Base:

Path elements

Query parameters

FeedNa:

siteld:

languag:

output:

satitle:

from:

submits:

**Fetch**

URL

**Priced Above (enter 2)**

Name:

Prompt:

Position:

Default:

Debug:

**Priced Below (enter 2)**

Name:

Prompt:

Position:

Default:

Debug:

**About this Pipe**

**eBay Price Watch**

This pipe is designed to use eBay's RSS API to find items within a certain price range.

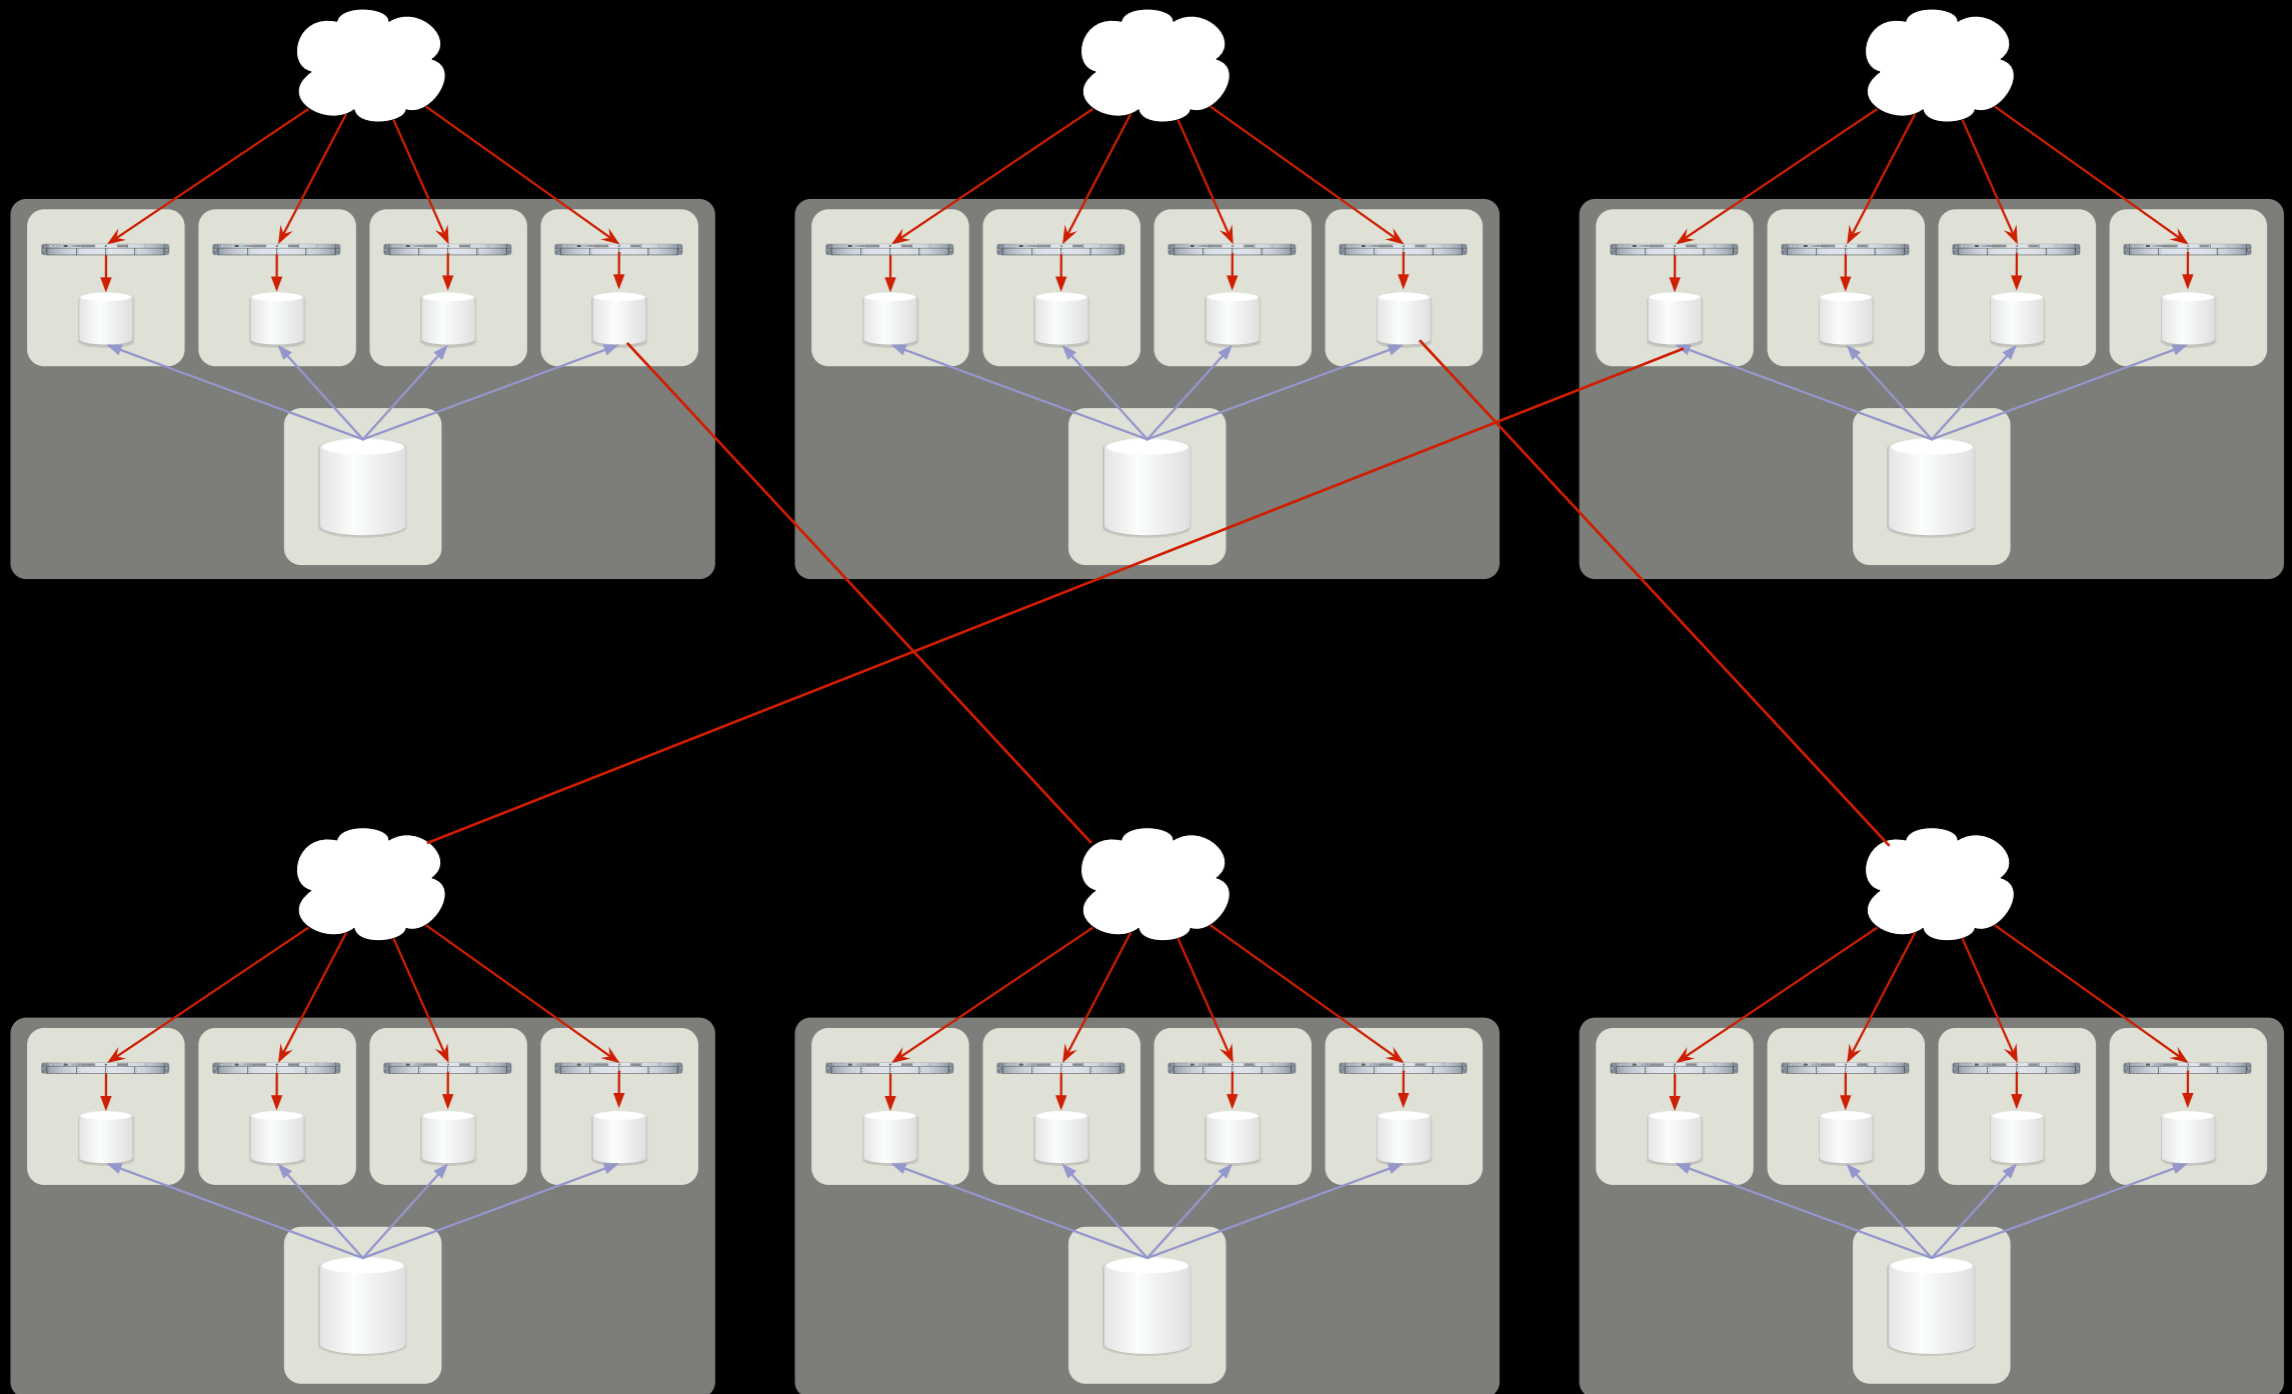
To do this we use the URLBuilder module to create a query. You can cut and paste REST urls into the "Base" field of URLBuilder and it will extract the key value pairs for you. We then wire a TextInput to the 'satitle' field so that the user's input will be supplied as the query term to eBay.

eBay supplements their RSS responses with a number of extensions. All of these are available to you to use in other modules. To filter on price, we use a Filter module and add two rules. In this example, we use 2 TextInputs to allow the user to set a minimum and maximum price range to filter the results.

You can make a copy of this Pipe right here to play with or try [running the pipe](#).

OK

# Building New Sites





# Requirements

# Requirements

massive scalability  
flexible deployment  
highly dynamic  
separation of concerns  
mashability

i.e.,

Services

Scalability

Simplicity


Reuse

Interoperability


i.e.,

**HTTP**


**Front-End**




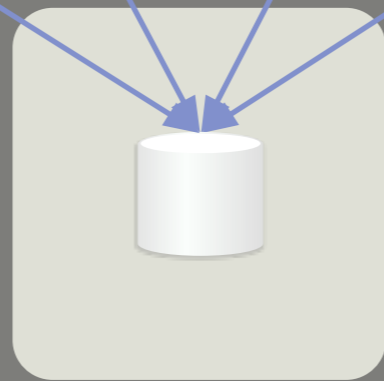
**Caching**




**Back-End**




**Database**





Single Source of Truth




Single Source of Truth

Cache Replicates Naturally



Adding Capacity Is Easy

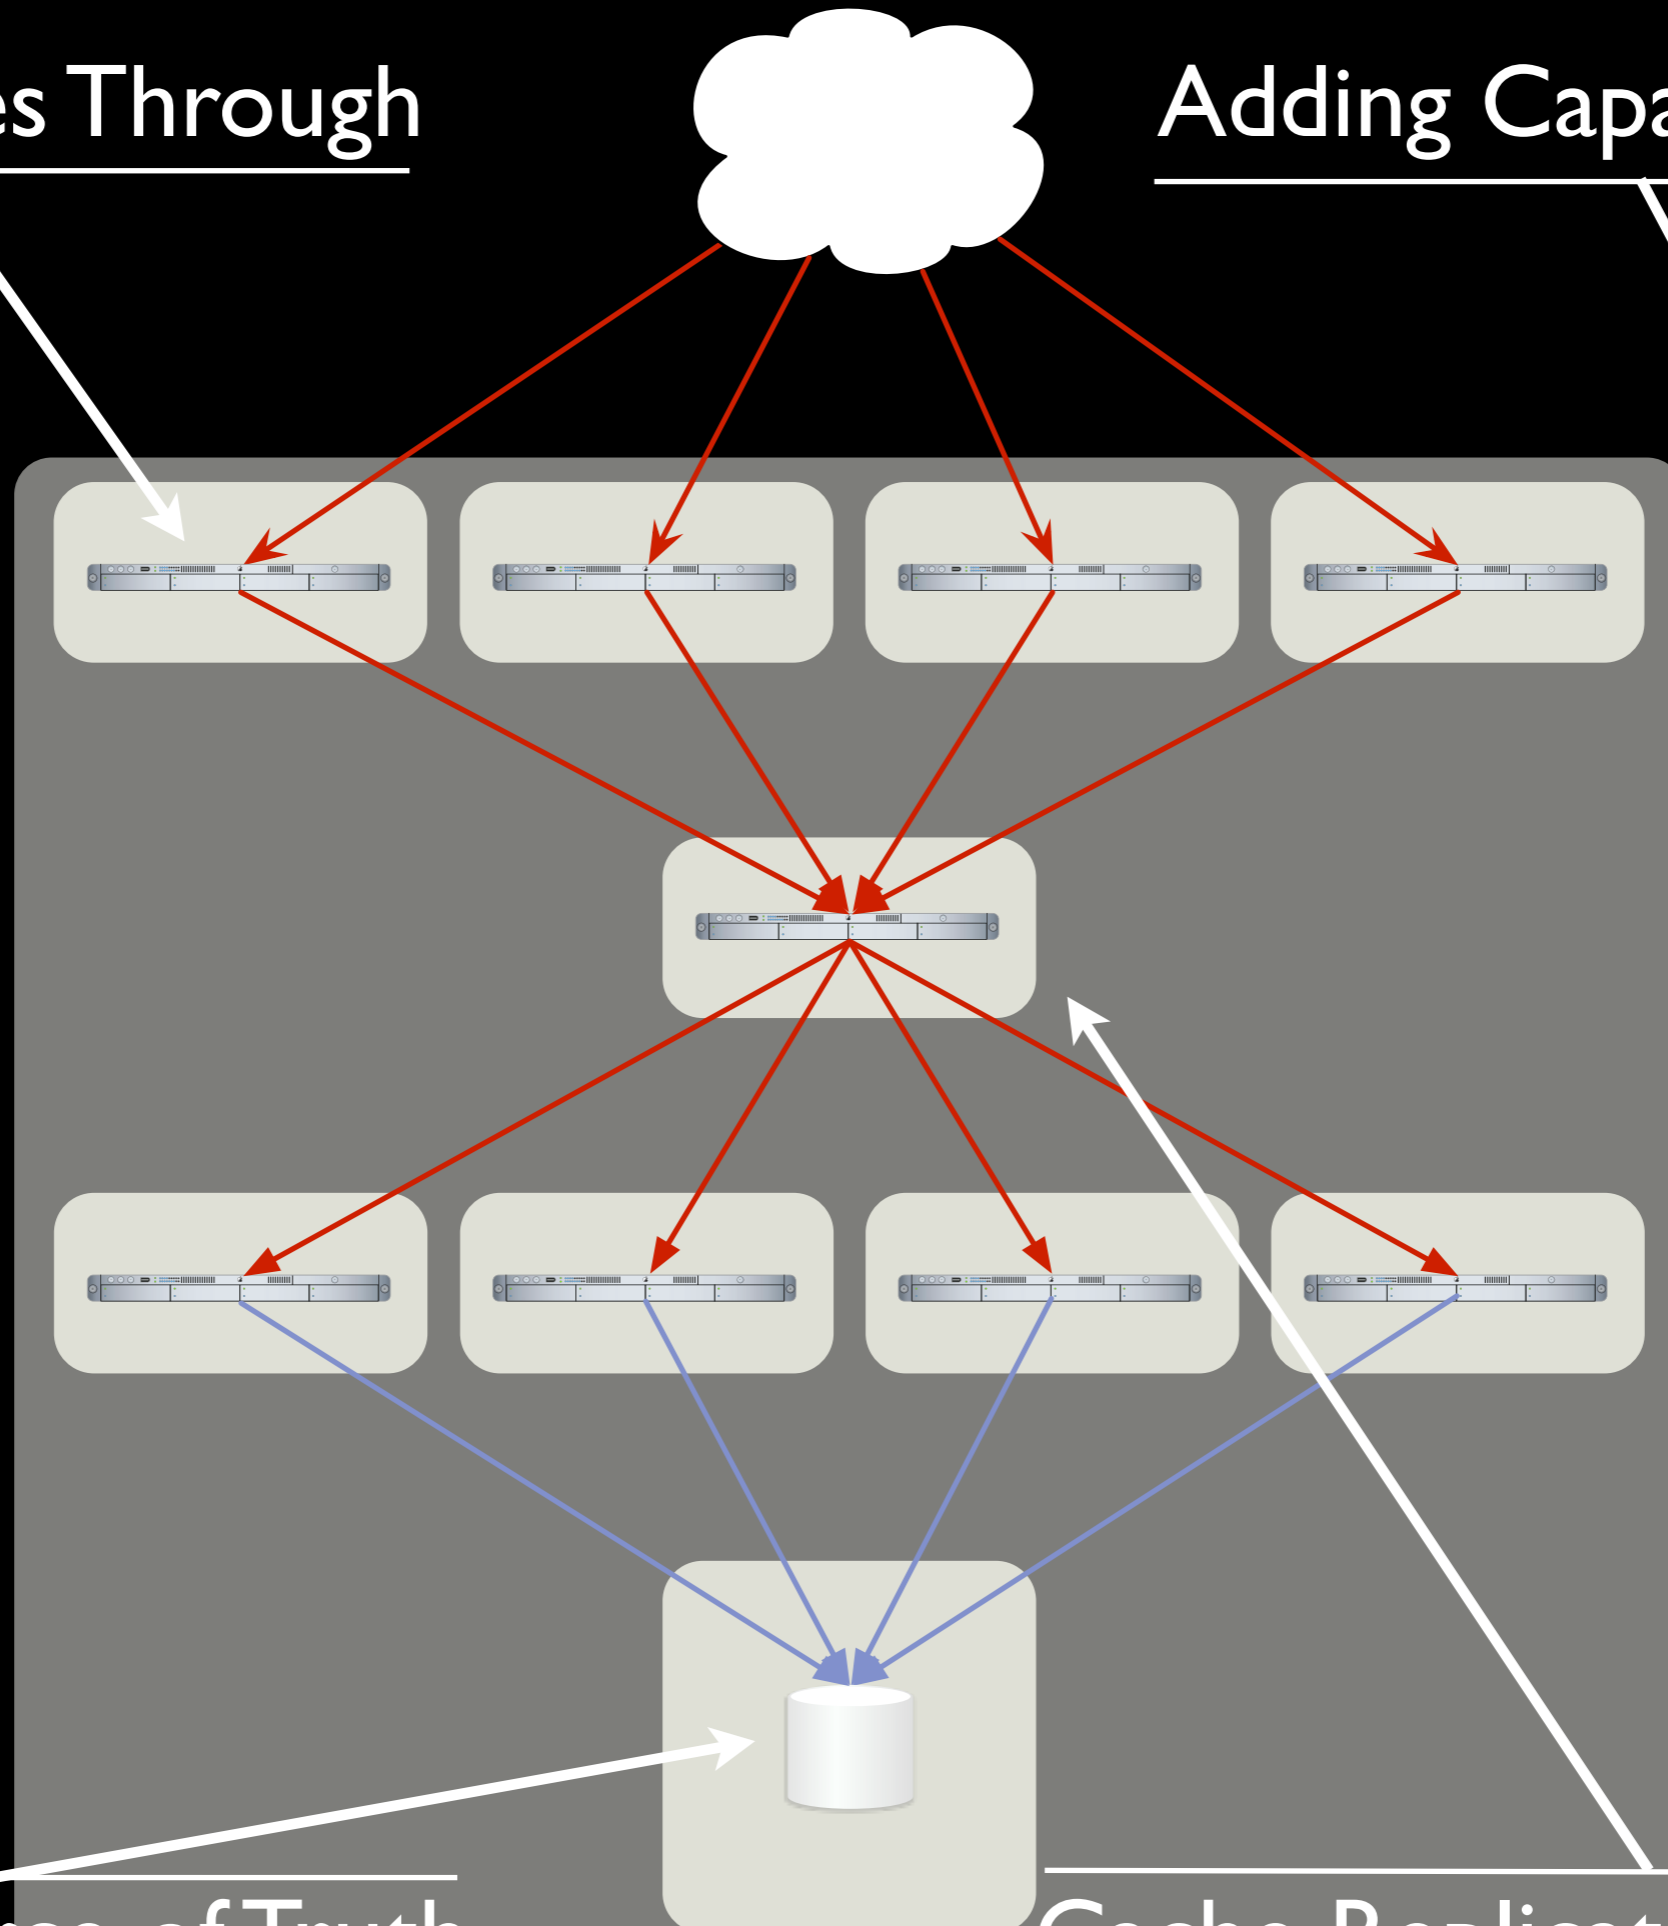


Single Source of Truth

Cache Replicates Naturally

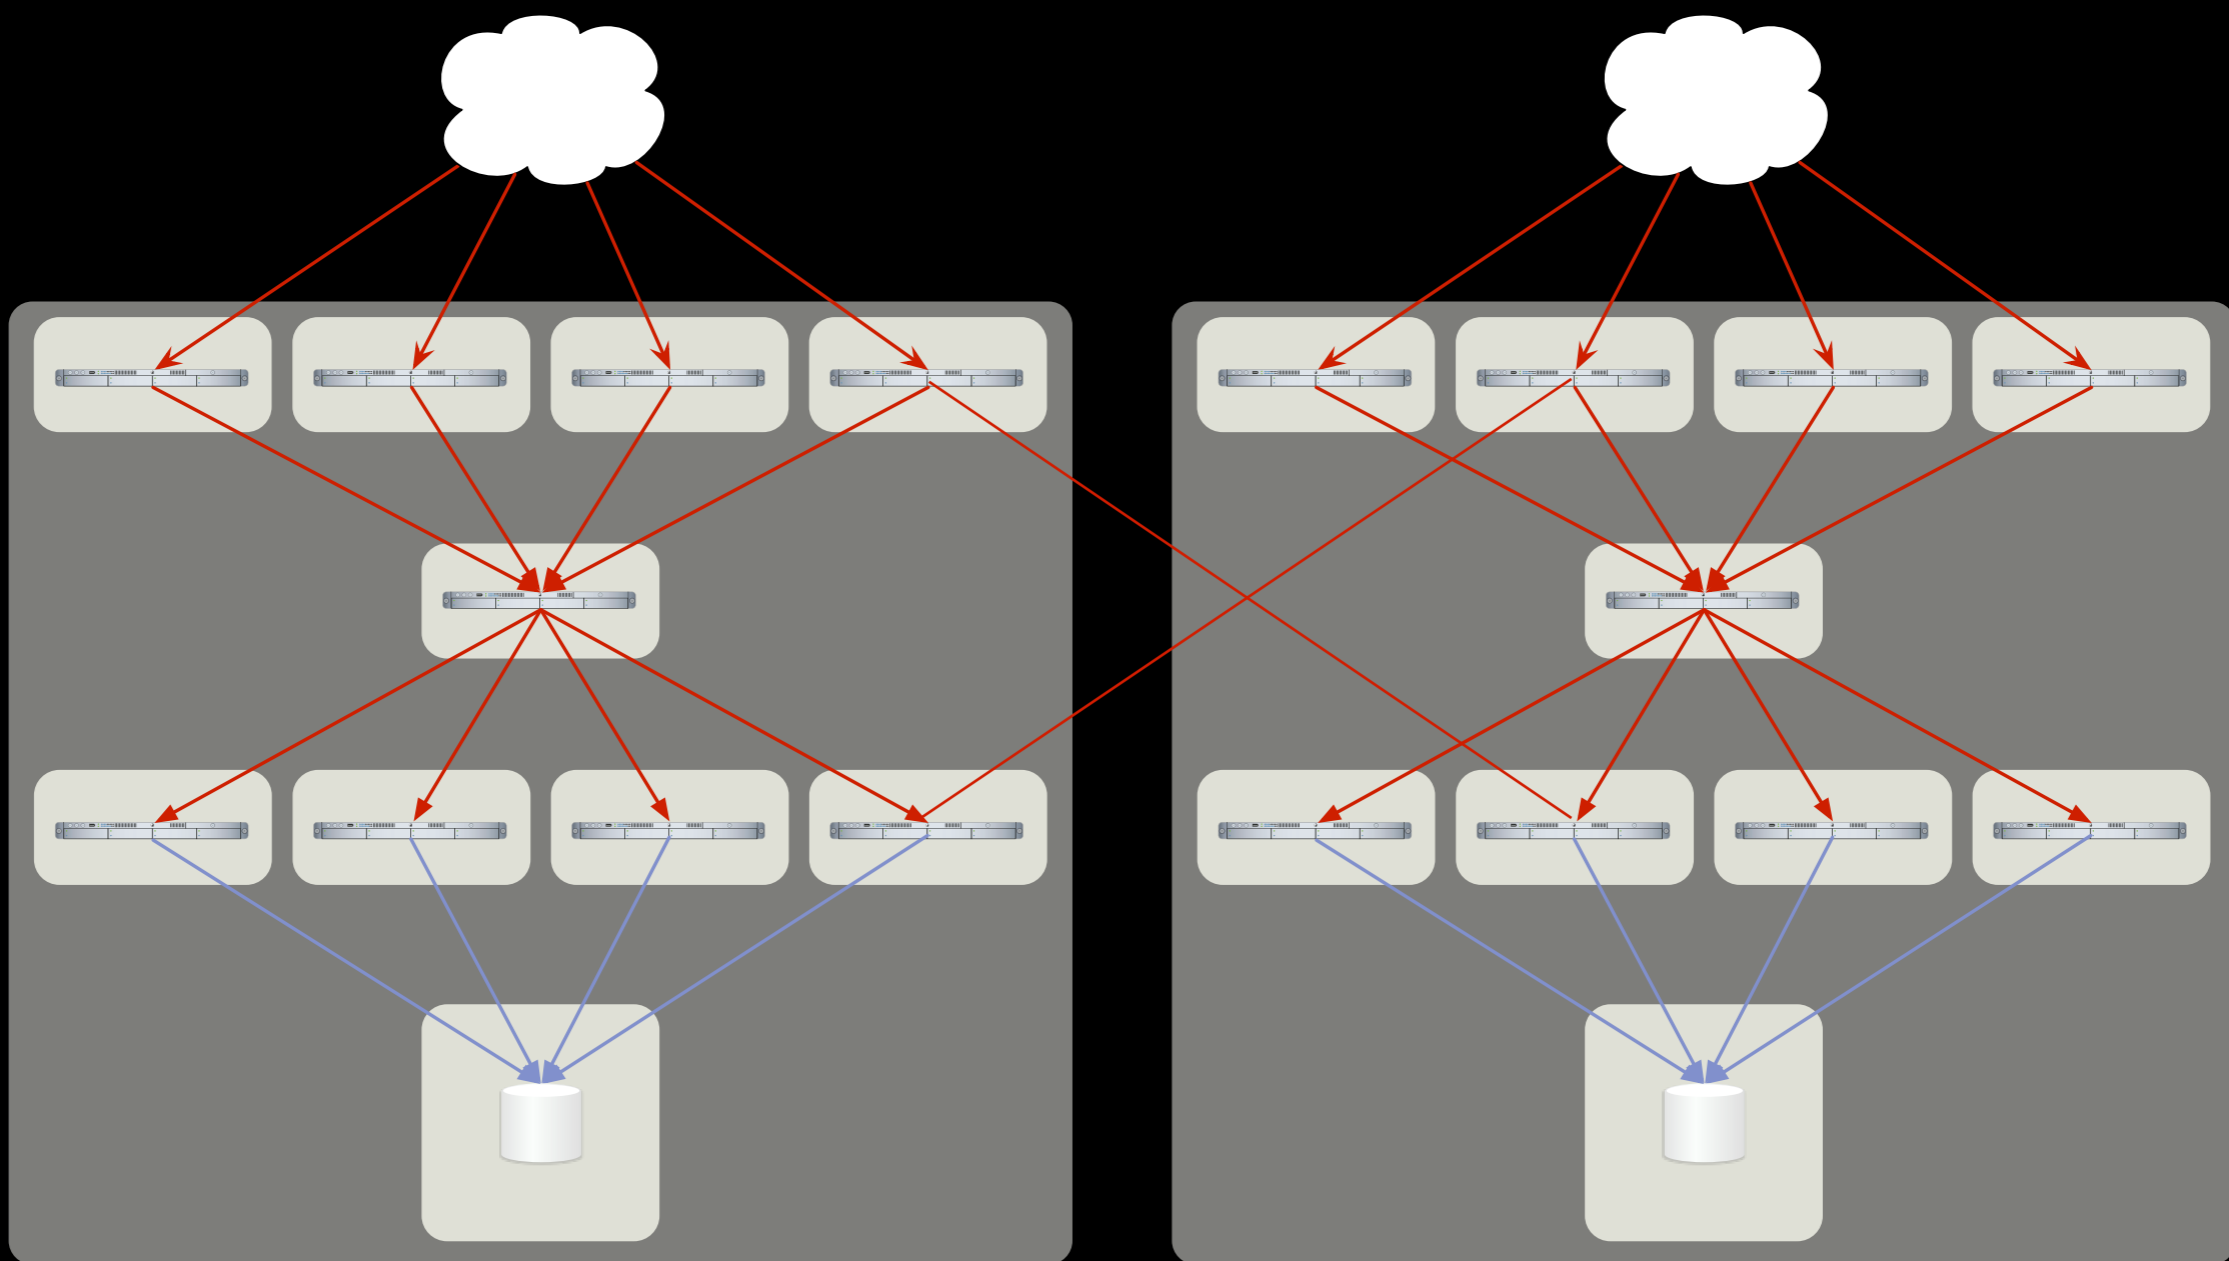
UGC Pushes Through

Adding Capacity Is Easy




Single Source of Truth

Cache Replicates Naturally




# HTTP Caching Intermediaries



# HTTP Caching Intermediaries


Freshness



# HTTP Caching Intermediaries

Freshness

Validation




# HTTP Caching Intermediaries

Freshness

Validation

Metrics



0101011010010101101010011001010101

# Service response times by cache result code

- 100019 log lines analysed, 0 parsing errors
- Start: Fri Sep 8 20:37:49 2006
- End: Fri Sep 8 21:14:30 2006
- 6445 distinct URLs seen, showing top 100

Response time histograms are a one-second window. See also [Squid's result code documentation](#).

url	accesses	CLIENT_REFRESH_MISS	HIT	MEM_HIT	MISS	NEGATIVE_HIT	SWAP
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	13360	0%	0%	84%	10%	6%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	7935	0%	0%	40%	60%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	7905	0%	0%	98%	1%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	7171	0%	1%	73%	25%	0%	1%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	6849	1%	0%	80%	18%	0%	1%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	4201	0%	0%	100%	0%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	3968	0%	0%	0%	93%	7%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	3957	0%	0%	100%	0%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	3069	0%	0%	100%	0%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1505	0%	0%	90%	10%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1386	0%	0%	82%	17%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1236	0%	0%	0%	100%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1189	0%	0%	96%	4%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1039	0%	0%	94%	5%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1028	0%	0%	95%	5%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	1018	1%	0%	75%	24%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	996	0%	0%	0%	100%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	799	1%	0%	66%	34%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	787	1%	0%	75%	24%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	699	0%	0%	99%	1%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	684	0%	0%	100%	0%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	658	0%	0%	100%	0%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	551	0%	0%	0%	100%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	524	0%	0%	98%	2%	0%	0%
http://twiki.corp.yahoo.com/pub/Platform/SquidLogMetrics/test.html	503	28%	0%	54%	17%	0%	0%
url	accesses	CLIENT_REFRESH_MISS	HIT	MEM_HIT	MISS	NEGATIVE_HIT	SWAP




# HTTP Caching Intermediaries

Freshness

Validation

Metrics

Load Balancing

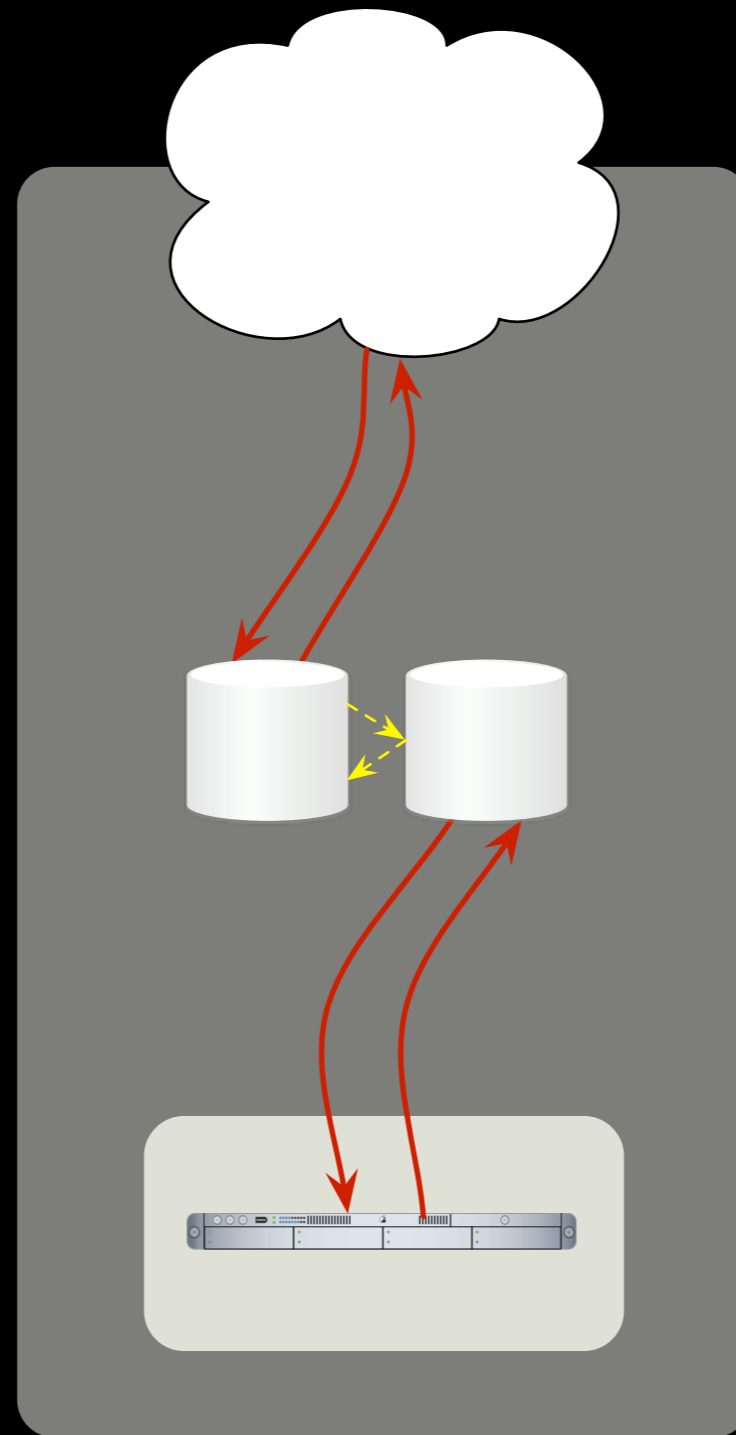


# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing


Cache Peering

# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing

Cache Peering


Negative Caching

# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing

Cache Peering

Negative Caching


Collapsed Forwarding

# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing

Cache Peering

Negative Caching

Collapsed Forwarding


stale-while-revalidate

# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing

Cache Peering

Negative Caching

Collapsed Forwarding

stale-while-revalidate

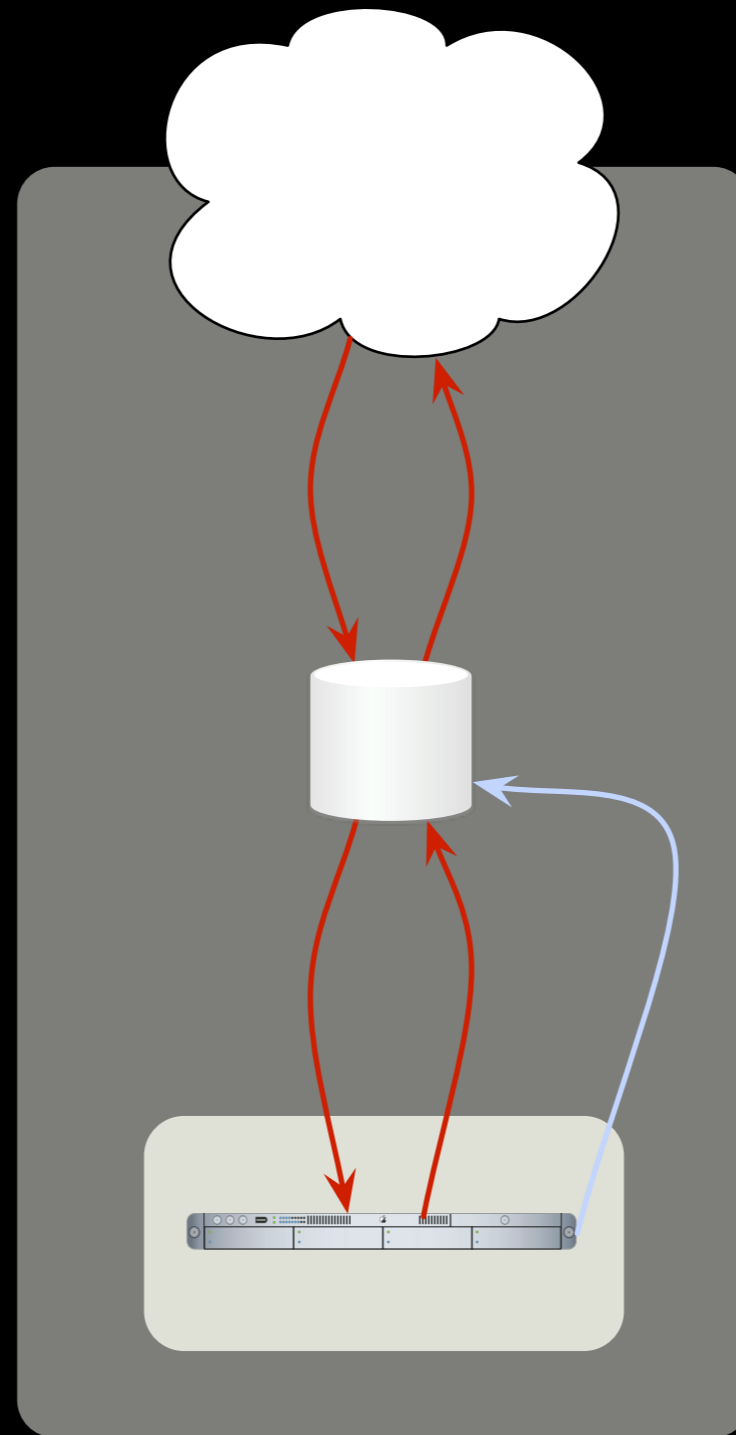
stale-if-error

# HTTP Caching Intermediaries

Freshness

Validation

Metrics



Load Balancing

Cache Peering

Negative Caching



Collapsed Forwarding

stale-while-revalidate

stale-if-error

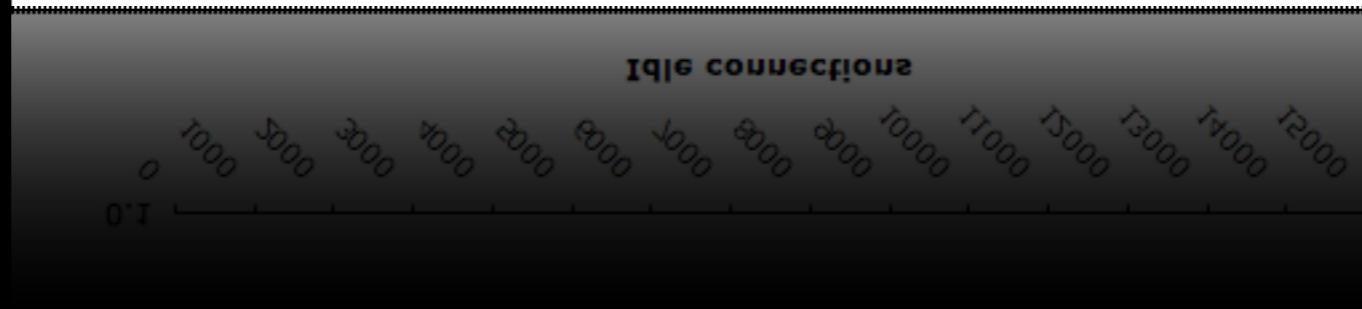

Invalidation Channels

# 12,000 req/sec/core





# Tens of Thousands of Connections



*Squid*

*Squid*

# pitfalls

REST vs. WS-\* wars

theory vs. practice

human-intuitive, but not programmer-intuitive

different deployment/operational concerns

formats are hard

format / interface proliferation

authentication isn't there yet

tools have a way to go

# what's needed

tools

web description language

data-oriented schema language

investment in the Atom stack

HTTP test suite

